

An Improved Filtering Algorithm for Big Read Datasets

Axel Wedemeyer^{*1}, Lasse Kliemann^{†1}, Anand Srivastav¹, Christian Schielke¹,
Thorsten B. Reusch², and Philip Rosenstiel³

¹Department of Computer Science, Kiel University, Christian-Albrechts-Platz 4,
24118 Kiel, Germany

²Marine Ecology, GEOMAR Helmholtz Centre for Ocean Research Kiel,
Düsternbrooker Weg 20, 24105 Kiel, Germany

³Institute of Clinical Molecular Biology, Kiel University, Schittenhelmstr. 12, 24105 Kiel,
Germany

October 12, 2016

Abstract

For single-cell or metagenomic sequencing projects, it is necessary to sequence with a very high mean coverage in order to make sure that all parts of the sample DNA get covered by the reads produced. This leads to huge datasets with lots of redundant data. A filtering of this data prior to assembly is advisable. Titus Brown et al. (2012) presented the algorithm Diginorm for this purpose, which filters reads based on the abundance of their k -mers. We present Bignorm, a faster and quality-conscious read filtering algorithm. An important new feature is the use of phred quality scores together with a detailed analysis of the k -mer counts to decide which reads to keep. With recommended parameters, in terms of median we remove 97.15% of the reads while keeping the mean phred score of the filtered dataset high. Using the SDAdes assembler, we produce assemblies of high quality from these filtered datasets in a fraction of the time needed for an assembly from the datasets filtered with Diginorm. We conclude that read filtering is a practical method for reducing read data and for speeding up the assembly process. Our Bignorm algorithm allows assemblies of competitive quality in comparison to Diginorm, while being much faster. Bignorm is available for download at <https://git.informatik.uni-kiel.de/axw/Bignorm.git>

^{*}axw@informatik.uni-kiel.de

[†]lki@informatik.uni-kiel.de

1 Background

Next generation sequencing systems (such as the Illumina platform) tend to produce an enormous amount of data — especially when used for single-cell or metagenomic protocols — of which only a small fraction is essential for the assembly of the genome. It is thus advisable to filter that data prior to assembly.

1.1 Problem Formulation

In order to describe our algorithm and its comparison, we need some formal definitions and concepts. Denote $\mathbb{N} := \{0, 1, 2, \dots\}$ the set of non-negative integers, and for each $n \in \mathbb{N}$ denote $[n] := \{1, \dots, n\}$ the integers from 1 to n (including 1 and n). Denote $\Sigma := \{A, C, G, T, N\}$ the alphabet of nucleotides plus the symbol N used to indicate an undetermined base. By Σ^* we denote all the finite strings over Σ , and for a $k \in \mathbb{N}$ by Σ^k all the strings over Σ of exactly length k . For $v \in \Sigma^*$, denote $|v| \in \mathbb{N}$ its length and $\bar{v} \in \Sigma^*$ its reverse complement. For $v, w \in \Sigma^*$, we write $v \cong w$ if $|v| = |w|$ and the two strings are equal up to places where either of them has the N symbol.

The input to the filter algorithm is a *dataset* $D = (n, m, R, Q)$ where for each $i \in [n]$ we have:

- $m(i) \in \mathbb{N}$: a flag for an unpaired ($m(i) = 1$) or paired ($m(i) = 2$) dataset;
- $R(i, s) \in \Sigma^*$ for each $s \in [m(i)]$: the set of *reads* in the dataset;
- $Q(i, s) \in \mathbb{Z}^{|R(i, s)|}$ for each $s \in [m(i)]$: the set of corresponding *phred scores*.

Each read $i \in [n]$ consists of $m(i)$ *read strings* $R(i, 1), \dots, R(i, m(i))$. For $t \in [|R(i, s)|] = \{1, \dots, |R(i, s)|\}$ we denote the nucleotide at position t in read string $R(i, s)$ by $R_t(i, s)$ and its phred score by $Q_t(i, s)$. Note that in terms of read strings, D may contain the “same” read multiple times (perhaps with different quality values), that is, there can be $i \neq j$ such that $R(i) = R(j)$. Hence it is beneficial that we refer to reads by their indices $1, \dots, n$.

Denote $x \in \Sigma^*$ the genome from which the reads were obtained and $g := |x|$ its length. (For the purpose of this exposition, we simplify by assuming the genome is a single string.). For each locus $\ell \in [g]$, the *coverage* $c_\ell(D)$ of ℓ with respect to D is informally described as the number of read strings that were or could have been produced by the sequencing machine while reading a part of x that contains locus ℓ . More precisely, for each $v \in \Sigma^*$ define

- $\triangleright c_\ell(v) := 1$ if there is a substring w of x which contains locus ℓ and satisfies $v \cong w$ or $v \cong \bar{w}$;
- $\triangleright c_\ell(v) := 0$ otherwise.

Then we define:

$$c_\ell(D) := \sum_{i=1}^n \sum_{s=1}^{m(i)} c_\ell(R(i, s))$$

A coverage of $c_\ell(D) \approx 20$ for each $\ell \in [g]$ has been empirically determined as optimal for a successful assembly of x from D [37]. On the other hand, in many setups, the coverage for a large number of loci is much higher than 20, often rising up to tens or hundreds of thousands, especially for single-cell or metagenomic protocols (see Table 3, “max” column for the maximal coverage of the datasets that we use in our experiments). In order to speed up the assembly process — or in extreme cases to make it possible in the first place, given certain restrictions on available RAM and/or time — a sub-dataset $D' = (n', m', R', Q')$ of D should be determined such that n' is much smaller than n while not losing essential information. The goal is that using D' , an assembly of similar quality than using D is possible. We only consider the natural approach to create D' by making a choice for each $i \in [n]$ whether to include read i in D' or not, so in particular $(R'(1), \dots, R'(n'))$ will be a sub-vector of $(R(1), \dots, R(n))$. When we include a read in D' , we also say that it is *accepted*, whereas when we exclude it, we say it is *rejected*. On an abstract level, a filtered dataset based on D can be specified by giving a set of indices $A \subseteq [n]$ that consists of exactly the accepted reads.

Many popular assemblers, such as SPAdes [16], Platanus [27], or Allpaths-LG [24], work with the *de Bruijn graph*, that is based on k -mers. Fix a parameter $k \in \mathbb{N}$; typically $21 \leq k$. The set of k -mers of a string $v \in \Sigma^*$, denoted $M(v, k) \subseteq \Sigma^k$, is the set of all strings of length k that are substrings of v . Sometimes we need to consider a k -mer multiple times if it occurs in multiple places in the string, and the corresponding set is denoted:

$$M^*(v, k) := \left\{ (\mu, p) \in \Sigma^k \times \mathbb{N} ; \mu \text{ is a substring of } v \text{ starting at position } p \right\}$$

For a read $i \in [n]$ and read string $s \in [m(i)]$ define $M(i, s, k) := M(R(i, s), k)$ and $M(i, k) := \bigcup_{s=1}^{m(i)} M(i, s, k)$, so $M(i, k)$ are all the k -mers that occur in any of the $m(i)$ read strings of $R(i)$. Denote also $M^*(i, s, k) := M^*(R(i, s), k)$.

1.2 Previous Work

We briefly survey two prior approaches for read pre-processing, namely *trimming* and *error correction*. Read trimming programmes (see [21] for a recent review) try to cut away the low quality parts of a read (or drop reads whose overall quality is low). These algorithms can be classified in two groups: *running sum* (Cutadapt, ERNE, SolexaQA with `-bwa` option) [19, 31, 32] and *window based* (ConDeTri, FASTX, PRINSEQ, Sickle, SolexaQA, and Trimmomatic) [12, 17, 19, 26, 35, 36]. The running sum algorithms take a quality threshold Q as input, which is subtracted from the phred score of each base of the read. The algorithms vary in the functions applied to the differences to determine the quality of a read, the direction in which the read is processed, the function’s quality threshold upon which the cutoff point is determined, and the minimum length of a read after the cutoff to be accepted.

The window based algorithms on the other hand first cut away the reads's 3' or 5' ends (depending on the algorithm) whose quality is below a specified minimum quality parameter and then determine a contiguous sequence of high quality using techniques similar to those used in the running sum algorithms.

All of these trimming algorithms generally work on a per-read basis, reading the input once and processing only a single read at a time. The drawback of this approach is that low quality sequences within a read are being dropped even when these sequences are not covered by any other reads whose quality is high. Also the phred score of a base is not independent between reads, i. e., a base whose phred score is low in one read is likely to have a low phred score in other reads as well and thus this low quality segment might get dropped altogether, creating uncovered regions. On the other hand sequences whose quality and abundance are high are added over and over although their coverage is already high enough, which yields higher memory usage than necessary.

Most of the error correction programs (see [15] for a recent review) read the input twice: a first pass gathers statistics about the data (often k -mer counts) which in a second pass are used to identify and correct errors. Some programs trim reads which cannot be corrected. Again, coverage is not a concern: reads which seem to be correct or which can be corrected are always accepted. According to [15], the probably best known and most used error correction program is Quake [29]. Its algorithm is based on two assumptions:

- “For sufficiently large k , almost all single-base errors alter k -mers overlapping the error to versions that do not exist in the genome. Therefore, k -mers with low coverage, particularly those occurring just once or twice, usually represent sequencing errors.”
- Errors follow a Gamma distribution, whereas true k -mers are distributed as per a combination of the Normal and the Zeta distribution.

In the first pass of the program, a score (based on the phred quality scores of the individual nucleotides) is computed for each k -mer. After this, Quake computes a *coverage cutoff* value, that is, the local minimum of the k -mer spectrum between the Gamma and the Normal maxima. All k -mers having a score higher than the coverage cutoff are considered to be correct (*trusted* or *solid* in error correction terminology), the others are assumed to be erroneous. In a second pass, Quake reads the input again and tries to replace erroneous k -mers by trusted ones using a maximum likelihood approach. Reads which cannot be corrected are optionally trimmed or dumped.

But the main goal of error correctors is not the reduction of the data volume (in particular, they do not pay attention to excessive coverage), hence they cannot replace the following approaches.

Titus Brown *et al.* invented an algorithm named *Diginorm* [37, 39] for read filtering that rejects or accepts reads based on the abundance of their k -mers. The name *Diginorm* is a short form for *digital normalization*: the goal is to normalize the coverage over all loci, using a computer

algorithm after sequencing. The idea is to reject those reads which mainly bring k -mers that have been seen many times in other reads already. Diginorm processes reads one by one. Let the read currently processed be $i \in [n]$. For each k -mer $\mu \in \Sigma^k$, define

$$c(\mu, i) := \left| \left\{ j \in \mathbb{N} ; (j < i) \text{ and } (\mu \in M(j, k)) \text{ and (read } j \text{ was previously accepted)} \right\} \right| ,$$

which says in how many accepted reads we have seen the k -mer μ so far. In order to save RAM, Diginorm does not keep track of those numbers exactly, but instead keeps appropriate estimates $\hat{c}(\mu, i)$ using the count-min sketch (CMS) [18]. For each $i \in [n]$ and $s \in [m(i)]$ denote the vector $C(i, s) := (\hat{c}(\mu, i))_{\mu \in M(i, s, k)}$. The read i is accepted if the median of the numbers in $C(i, s)$ is below a fixed threshold, usually 20, for each $s \in [m(i)]$. It was demonstrated that successful assemblies are still possible after Diginorm removed the majority of the data.

1.3 Our Algorithm

Diginorm is a pioneering work. However, the following points, which are important from the biological or computational quality point of view, are not covered in Diginorm. We present them as an enhancement in our work:

- (i) We incorporate the important phred quality score into the decision whether to accept or to reject a read, using a quality threshold. This allows a tuning of the filtering process towards high-quality assemblies, by using different thresholds.
- (ii) When deciding whether to accept or to reject read i , we do a detailed analysis of the numbers in the vectors $C(i, s)$. Diginorm merely considers their medians.
- (iii) We offer a better handling of the **N** case, that is, when the sequencing machine could not decide for a particular nucleotide. Diginorm simply converts all **N** to **A**, which can lead to false k -mer counts.¹
- (iv) We provide a substantially faster implementation. For example, we include fast hashing functions (see [22, 38]) for counting k -mers through the count-min sketch data structure (CMS), and we use the C programming language and OpenMP.

A detailed description of our algorithm, called *Bignorm*, is given in the next section. Its name was chosen to emphasize the goal of drastically reducing massive datasets.

Bignorm, like Diginorm, is based on the count-min sketch (CMS) for counting k -mers. CMS is a probabilistic data structure for counting objects from a large universe. We give a brief and abstract description. Let $a = (a_1, \dots, a_N) \in \mathbb{N}^N$ be a vector, given implicitly as a sequence of updates of the form (p, Δ) with $p \in [N]$ and $\Delta \in \mathbb{N}$. Each update (p, Δ) modifies a in the way $a_p := a_p + \Delta$; where initially $a = (0, \dots, 0)$. If $\Delta = 1$ in each update, then an interpretation

¹We have observed some evidence that this may lead to a spuriously higher GC content. This will be investigated in future work.

of the vector a is that we count how many times we observe each of the objects identified by the numbers in $[N]$. If N is large, e. g., if N is the number 4^k of all possible k -mers (we do not count k -mers with \mathbf{N} symbols), then we may not be able to store a in RAM. (For example, the typical choice of $k = 21$ brings a into terabyte range; in our experiments we use $k = 32$.) Instead we fix two parameters: the *width* $m \in \mathbb{N}$ and the *depth* $t \in \mathbb{N}$ and store a matrix of $m \cdot t$ CMS counters $c_{p,q}$ with $p \in [m]$ and $q \in [t]$. Moreover, we randomly draw t hash functions h_1, \dots, h_t from a universal family. Each h_q maps from $[N]$ to $[m]$. Initially, all counters in the matrix are zero. Upon arrival of an update (p, Δ) , for each row $q \in [t]$ we update $c_{h_q(p),q} := c_{h_q(p),q} + \Delta$. That is, for each row q we use the hash function h_q to map from the larger space $[N]$ (from which the index p comes) to the smaller space $[m]$ of possible positions in the row. Denote

$$\hat{a}_p := \min \{c_{h_1(p),1}, \dots, c_{h_t(p),t}\} . \quad (1)$$

Then it can be proved [18] that \hat{a}_p is an estimate of a_p in the following sense: clearly $a_p \leq \hat{a}_p$, and with probability at least $1 - e^{1-t}$ we have $\hat{a}_p \leq \frac{e}{m-1} \sum_{j=1}^N a_j$. The probability is over the choice of hash functions. For example, choosing $t := 10$ is enough to push the error probability, upper-bounded by e^{1-t} , below 0.013%.

In our application, $N = 4^k$ is the number of possible k -mers (without \mathbf{N} symbols) and we implement a bijection $\beta : \Sigma^k \rightarrow [N]$, so we can identify each k -mer μ by a number $\beta(\mu) \in [N]$. Upon accepting some read i , we update the CMS counters using all the updates of the form $(\beta(\mu), 1)$ with $\mu \in M(i, k)$ not containing the \mathbf{N} symbol, that is, for each such μ we increase the count $\beta(\mu)$ by $\Delta = 1$. Then when all the reads $1, \dots, i-1$ have been processed, the required count $c(\mu, i)$ corresponds to the entry $a_{\beta(\mu)}$ in the vector a used in the description of CMS, and for the estimate $\hat{c}(\mu, i)$ we can use the estimate $\hat{a}_{\beta(\mu)}$ as given in (1).

2 Methods

2.1 Description of Bignorm

We give a detailed description of our enhancements (i) to (iv) that were briefly lined out on the preceding page. Although most of the settings are generic, in some places we assume that data comes from the Illumina.

We start with (i), (ii), and (iii). Fix a read $i \in [n]$ and a read string $s \in [m(i)]$. Recall that for each $t \in [|R(i, s)|]$ the nucleotide $R_t(i, s)$ at position t in the read string $R(i, s)$ is associated with a quality value $Q_t(i, s)$ known as *phred score*. We want to assign a single value $Q(i, s, \mu, p)$ to each $(\mu, p) \in M^*(i, s, k)$. We do so by taking the minimum phred score over the nucleotides in μ when aligned at position p , that is:

$$Q(i, s, \mu, p) := \min_{t=p}^{p+k-1} Q_t(i, s)$$

(μ occurs on the right-hand side only implicitly through its length k .)

Fix the following parameters:

- *N-count threshold* $N_0 \in \mathbb{N}$, which is 10 by default;
- *quality threshold* $Q_0 \in \mathbb{Z}$, which is 20 by default;
- *rarity threshold* $c_0 \in \mathbb{N}$, which is 3 by default;
- *abundance threshold* $c_1 \in \mathbb{N}$, which is 20 by default;
- *contribution threshold* $B \in \mathbb{N}$, which is 3 by default.

When our algorithm has to decide whether to accept or to reject a read $i \in [n]$, it performs the following steps. If the number of **N** symbols counted over all $m(i)$ read strings in i is larger than N_0 , the read is rejected right away. Otherwise, for each $s \in [m(i)]$ define the set of *high-quality k-mers*:

$$H(s) := \left\{ (\mu, p) \in M^*(i, s, k) ; (Q_0 \leq Q(i, s, \mu, p)) \text{ and } (\mu \text{ does not contain } \mathbf{N}) \right\}$$

We determine the *contribution* of $R(i, s)$ to k -mers of different frequencies:

$$\begin{aligned} b_0(s) &:= |\{(\mu, p) \in H(s) ; \widehat{c}(\mu, i) < c_0\}| \\ b_1(s) &:= |\{(\mu, p) \in H(s) ; c_0 \leq \widehat{c}(\mu, i) < c_1\}| \end{aligned}$$

Note that the frequencies are determined via CMS counters and do not consider the position p at which the k -mer is found in the read string. The read i is accepted if and only if at least one of the following conditions is met:

$$b_0(s) > k \text{ for at least one read string } s \tag{2}$$

$$\sum_{s=1}^{m(i)} b_1(s) \geq B \tag{3}$$

If the read is accepted, then for each $\mu \in M(i, k)$ the corresponding CMS counter is incremented, provided that μ does not contain the **N** symbol. Then processing of the next read starts.

The rationale for condition (2) is as follows. If a k -mer is seen less than c_0 times, we suspect it to be the result of a read error. However, if more than k k -mers in a read string contain an error, this read string must have more than one erroneous nucleotide. This is not likely for the Illumina platform, since there, most errors are single substitutions [29]. So if $b_0(s) > k$ for some s , then the read string $R(i, s)$ should be assumed to correctly contain a *rare* k -mer, so it must not be filtered out.

Condition (3) says that in the read i , there are enough (namely at least B) k -mers where each of them is too frequent to be a read error (CMS counters at least c_0) but not so abundant that it should be considered redundant (CMS counters less than c_1).

This concludes the description of (i), (ii), and (iii), namely how we analyze the counts in $C(i, s) = (\hat{c}(\mu, i))_{\mu \in M(i, s, k)}$ for each read i and $s \in [m(i)]$, how we incorporate quality information, and how we handle the N symbol.

Finally, to accomplish (iv), we wrote a multi-threaded implementation completely in the C programming language. The parallel code uses OpenMP. For comparison, the implementation of the original Diginorm algorithm (included in the khmer-package [20]) features a single-threaded design and is written in Python and C++; strings have to be converted between Python and C++ at least twice.

2.2 Experimental Setup

For the experimental evaluation, we collected the following datasets. We use two single cell datasets of the UC San Diego, one of the group of Ute Hentschel (now GEOMAR Kiel) and 10 datasets from the JGI Genome Portal. The datasets from JGI were selected as follows. On the JGI Genome Portal [13], we used “single cell” as search term. We narrowed the results down to datasets which had all of the following properties:

- status “complete”;
- containing read data *and* an assembly in the download section;
- aligning the reads to the assembly using bowtie2 [30] yields an “overall alignment rate” of more than 70%.

From those datasets, we arbitrarily selected one per species, until we had a collection of 10 datasets. We refer to each combination of species and selected dataset as a *case* in the following. In total, we have 13 cases; the details are given in Table 1.

Short Name	Species/Description	Source	URL
ASZN2	Candidatus Poribacteria sp. WGA-4E_FD	Hentschel Group [28]	[7]
Aceto	Acetothermia bacterium JGI MDM2 LHC4sed-1-H19	JGI Genome Portal	[1]
Alphaproteo	Alphaproteobacteria bacterium SCGC AC-312_D23v2	JGI Genome Portal	[2]
Arco	Arcobacter sp. SCGC AAA036-D18	JGI Genome Portal	[3]
Arma	Armatimonadetes bacterium JGI 0000077-K19	JGI Genome Portal	[4]
Bacteroides	Bacteroidetes bacVI JGI MCM14ME016	JGI Genome Portal	[5]
Caldi	Calescamantes bacterium JGI MDM2 SSWTFF-3-M19	JGI Genome Portal	[6]
Cauro	Caulobacter bacterium JGI SC39-H11	JGI Genome Portal	[8]
Chloroflexi	Chloroflexi bacterium SCGC AAA257-O03	JGI Genome Portal	[9]
Crenarch	Crenarchaeota archaeon SCGC AAA261-F05	JGI Genome Portal	[10]
Cyanobact	Cyanobacteria bacterium SCGC JGI 014-E08	JGI Genome Portal	[11]
E.coli	E.coli K-12, strain MG1655, single cell MDA, Cell one	UC San Diego	[14]
SAR324	SAR324 (Deltaproteobacteria)	UC San Diego	[14]

Table 1. Selected Species and Datasets (Cases)

For each case, we analyze the results obtained with Diginorm and with Bignorm using quality parameters $Q_0 \in \{5, 8, 10, 12, 15, 18, 20, \dots, 45\}$. Analysis is done on the one hand in terms of data reduction, quality, and coverage. On the other hand, we study actual assemblies that are computed with SPAdes [16] based on the raw and filtered datasets. All the details are given in the next section.

The dimensions of the count-min sketch are fixed to $m = 1024$ and $t = 10$, thus 10 GB of RAM where used.

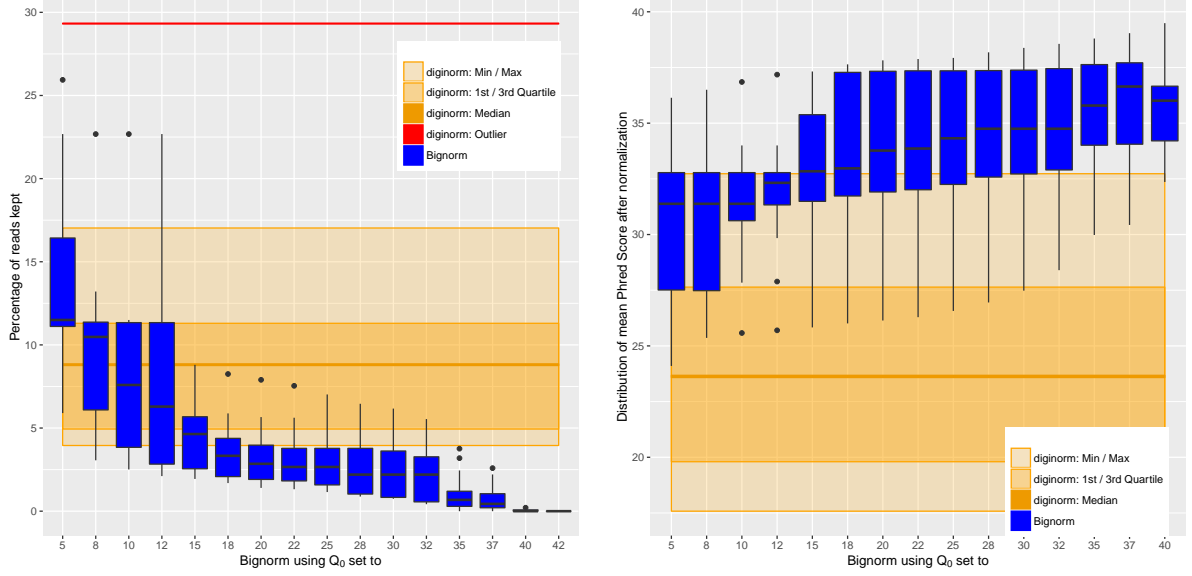
3 Results

We do analysis in large parts by looking at percentiles and quartiles. The i th quartile is denoted Q_i , where we use Q_0 for the minimum, Q_2 for the median, and Q_4 for the maximum. The i th percentile is denoted P_i ; we often use the 10th percentile P_{10} .

3.1 Number of Accepted Reads

Statistics for the number of accepted reads are given as a boxplot in Figure 1a on the next page. This plot is constructed as follows. Each of the blue boxes corresponds to Bignorm with a particular Q_0 , while Diginorm is represented as the wide orange box in the background (recall that Diginorm does not consider quality values). Note that the “whiskers” of Diginorm’s box are shown as light-orange areas. For each box, for each case the raw dataset is filtered using the algorithm and algorithmic parameters corresponding to that box, and the percentage of the accepted reads is taken into consideration. So for example, if the top of a box (which corresponds to the 3rd quartile, also denoted Q_3) gives the value $x\%$, then we know that for 75% of the cases, $x\%$ or less of the reads were accepted using the algorithm and algorithmic parameters corresponding to the box.

There are two prominent outliers: one for Diginorm with value $\approx 29\%$ (shown as the red line at the top) and one for Bignorm for $Q_0 = 5$ with value $\approx 26\%$. In both cases the Arma dataset is responsible, for which we do not have an explanation at this time. For $15 \leq Q_0$, even Bignorm’s outliers fall below Diginorm’s median, and for $18 \leq Q_0$ Bignorm keeps less than 5% of the reads for at least 75% of the datasets. In the range $20 \leq Q_0 \leq 25$, Bignorm delivers similar results for the different Q_0 , and the gain in reduction for larger Q_0 is small up to $Q_0 = 32$. For even larger Q_0 , there is another jump in reduction, but we will see that coverage and the quality of the assembly suffer too much in that range. We conjecture that in the range $18 \leq Q_0 \leq 32$, we remove most of the actual errors, whereas for larger Q_0 we also remove useful information.



(a) Percentage of accepted reads (i. e., reads kept) over all datasets. (b) Mean quality values of the accepted reads over all datasets.

Figure 1. Boxplots showing reduction and quality statistics.

3.2 Quality Values

Statistics for phred quality scores in the filtered datasets are given in Figure 1b on this page. The data was obtained using `fastx_quality_stats` from the FASTX Toolkit [12] on the filtered fastq files and calculating the mean phred quality scores over all read positions for each dataset. Looking at the statistics for these overall means, for $15 \leq Q_0$, Bignorm's median is better than Diginorm's maximum. For $20 \leq Q_0$, this effect becomes even stronger. For all Q_0 values, Bignorm's minimum is clearly above Diginorm's median. Note that 10 units more means reducing error probability by factor 10.

In Table 2, we give quartiles of mean quality values for the raw datasets and Bignorm's datasets produced with $Q_0 = 20$. Bignorm improves slightly on the raw dataset in all five quartiles.

Of course, all this could be explained by Bignorm simply cutting away any low-quality reads. However, the data in the next section suggests that Bignorm may in fact be more careful than this.

Quartile	Bignorm	raw
Q4 (max)	37.82	37.37
Q3	37.33	36.52
Q2 (median)	33.77	32.52
Q1	31.91	30.50
Q0 (min)	26.14	24.34

Table 2. Comparing quality values for the raw dataset and Bignorm with $Q_0 = 20$.

3.3 Coverage

In Figure 2 on page 13 we see statistics for the coverage. The data was obtained by remapping the filtered reads onto the assembly from the JGI using bowtie2 and then using `coverageBed` from the bedtools [33] and R [34] for the statistics. In Figure 2a, the mean is considered. For $15 \leq Q_0$, Bignorm reduces the coverage heavily. For $20 \leq Q_0$, Bignorm’s Q3 is below Diginorm’s Q1. This may raise the concern that Bignorm could create areas with insufficient coverage. However, in Figure 2b, we look at the 10th percentile (P10) of the coverage instead of the mean. We consider this statistics as an indicator for the impact of the filtering on areas with low coverage. For $Q_0 \leq 25$, Bignorm’s Q3 is on or above Diginorm’s maximum, and Bignorm’s minimum coincides with Diginorm’s (except for $Q_0 = 10$, where we are slightly below). In terms of median, both algorithms are very similar for $Q_0 \leq 25$. We consider all this as a strong indication that we cut away in the right places.

For $28 \leq Q_0$, there is a clear drop in coverage, so we do not recommend such Q_0 values.

In Table 3, we give coverage statistics for each dataset. The reduction compared to the raw dataset in terms of mean, P90, and maximum is substantial. But also the improvement of Bignorm over Diginorm in mean, P90, and maximum is considerable for most datasets.

3.4 Assessment through Assemblies

The quality and significance of read filtering is subject to complete assemblies, which is the final “road test” of algorithms. For each case, we do an assembly with SPAdes using the raw dataset and those filtered with Diginorm and Bignorm for a selection of Q_0 values. The assemblies are then analyzed using quast [25] and the assembly from the JGI as reference. Statistics for four cases are shown in Figure 3. We give the quality measures N50, genomic fraction, and largest contig, and in addition the overall running time (pre-processing plus assembly). Each measure is given in percent relative to the raw dataset.

Dataset	Algorithm	P10	mean	P90	max
Aceto	Bignorm	6	132	216	6801
	Diginorm	7	171	295	12020
	raw	15	9562	17227	551000
Alphaproteo	Bignorm	10	43	92	884
	Diginorm	7	173	481	6681
	raw	25	5302	14070	303200
Arco	Bignorm	1	98	54	2103
	Diginorm	1	362	200	6114
	raw	3	10850	4091	220600
Arma	Bignorm	8	23	32	358
	Diginorm	8	79	141	5000
	raw	17	629	1118	31260
ASZN2	Bignorm	40	70	83	2012
	Diginorm	23	143	354	3437
	raw	50	1738	4784	43840
Bacteroides	Bignorm	3	74	90	6768
	Diginorm	3	123	205	7933
	raw	7	6051	8127	570900
Caldi	Bignorm	25	63	110	786
	Diginorm	15	67	135	3584
	raw	27	1556	3643	33530
Caulo	Bignorm	7	228	216	10400
	Diginorm	8	362	491	35520
	raw	8	10220	9737	464300
Chloroflexi	Bignorm	8	72	101	2822
	Diginorm	9	412	878	20850
	raw	9	5612	7741	316900
Crenarch	Bignorm	8	104	159	3770
	Diginorm	10	560	1285	29720
	raw	10	8086	14987	316700
Cyanobact	Bignorm	9	144	153	5234
	Diginorm	10	756	1450	26980
	raw	10	9478	11076	356600
E.coli	Bignorm	37	45	56	234
	Diginorm	50	382	922	7864
	raw	112	2522	6378	56520
SAR324	Bignorm	24	49	71	1410
	Diginorm	18	53	107	2473
	raw	26	1086	2761	106000

Table 3. Coverage statistics for Bignorm with $Q_0 = 20$, Diginorm, and the raw datasets.




Figure 2. Boxplots showing coverage statistics.

Generally, our biggest improvements are for N50 and running time. For $15 \leq Q_0$, Bignorm is always faster than Diginorm, for three of the four cases by a large margin. In terms of N50, for $15 \leq Q_0$ we observe improvements for three cases. For E.coli, Diginorm's N50 is 100%, that we also attain for $Q_0 = 20$. In terms of genomic fraction and largest contig, we cannot always attain the same quality as Diginorm; the biggest deviation at $Q_0 = 20$ is 10 percentage points for the ASZN2 case. The N50 is generally accepted as one of the most important measures, as long as the assembly represents the genome well (as measured here by the genomic fraction) [23].

In Table 4, we give statistics for $Q_0 = 20$ and each case. In terms of genomic fraction, Bignorm is generally not as good as Diginorm. However, excluding the Aceto and Arco cases, Bignorm's genomic fraction is still always at least 95%. For Aceto and Arco, Bignorm misses 3.21% and 3.48%, respectively, of the genome in comparison to Diginorm. In 8 cases, Bignorm's N50 is better or at least as good as Diginorm's. The 4 cases where we have smaller N50 are Arco, Caldi, Caulo, Crenarch, and Cyanobact.

Bignorm's mean phred score is always slightly larger than that of the raw dataset, whereas Diginorm's is always smaller. For some cases, the difference is substantial; the quartiles for the ratio of Diginorm's mean phred score to that of the raw dataset are given in Table 6 in the first row.

Clearly, our biggest gain is in running time, for the filtering as well for the assembly. Quartiles of the corresponding improvements are given in rows two and three of Table 6.

Dataset	Algorithm	reads kept in %	mean phred score	contigs $\geq 10\,000$	filter time in sec	SPAdes time in sec
Aceto	Bignorm	3.16	37.33	1	906	1708
	Diginorm	3.95	27.28	1	3290	4363
	raw		36.52	3		47813
Alphaproteo	Bignorm	3.13	34.65	18	623	420
	Diginorm	7.81	28.73	17	1629	11844
	raw		33.64	17		29057
Arco	Bignorm	2.20	33.77	4	429	207
	Diginorm	8.76	21.39	6	1410	1385
	raw		32.27	6		15776
Arma	Bignorm	7.90	28.21	44	240	135
	Diginorm	29.30	21.19	50	588	1743
	raw		26.96	44		5371
ASZN2	Bignorm	5.66	37.66	118	1224	1537
	Diginorm	12.62	32.73	130	5125	21626
	raw		36.85	112		47859
Bacteroides	Bignorm	2.85	37.47	6	653	3217
	Diginorm	4.94	27.64	5	2124	3668
	raw		37.25	9		32409
Caldi	Bignorm	3.97	37.82	41	842	455
	Diginorm	5.61	30.67	36	1838	793
	raw		37.37	38		7563
Caulo	Bignorm	2.40	36.95	10	679	712
	Diginorm	4.70	25.16	9	2584	765
	raw		36.01	13		18497
Chloroflexi	Bignorm	1.40	31.91	32	694	134
	Diginorm	9.70	18.91	33	2304	1852
	raw		30.50	34		15108
Crenarch	Bignorm	1.46	33.18	19	1107	790
	Diginorm	9.72	19.80	18	2931	3754
	raw		31.49	26		20590
Cyanobact	Bignorm	1.65	30.45	12	679	450
	Diginorm	11.30	17.58	13	1487	1343
	raw		28.49	13		9417
E. coli	Bignorm	1.91	26.14	67	2279	598
	Diginorm	17.03	19.34	63	9105	3995
	raw		24.34	64		16706
SAR324	Bignorm	4.34	33.05	55	1222	708
	Diginorm	4.69	23.58	52	3706	3085
	raw		32.52	51		26237

Table 4. Filter and assembly statistics for Bignorm with $Q_0 = 20$, Diginorm and the raw datasets (I)

Dataset	Algorithm	N50			Longest Contig Length			Genomic Fraction			Misassembled Contig Length		
		% of abs	% of raw	% of Diginorm	% of abs	% of raw	% of Diginorm	% of abs	% of raw	% of Diginorm	% of abs	% of raw	% of Diginorm
Aceto	Bignorm	2324	79	105	11525	98	100	91	97	97	52487	148	178
	Diginorm	2216	76		11525	98		94	100		29539	84	
	raw	2935			11772			94			35351		
Alphaproteo	Bignorm	11750	94	115	43977	91	95	98	101	105	52001	120	89
	Diginorm	10213	82		46295	95		93	95		58184	134	
	raw	12446			48586			98			43388		
Arco	Bignorm	3320	81	97	12808	57	57	85	100	97	76797	99	91
	Diginorm	3434	84		22463	100		88	103		84613	109	
	raw	4092			22439			85			77888		
Arma	Bignorm	18432	102	107	108140	100	100	98	100	100	774291	91	103
	Diginorm	17288	96		108498	100		98	100		748560	88	
	raw	18039			108498			98			849085		
ASZN2	Bignorm	19788	91	88	72685	71	88	97	99	99	2753167	94	105
	Diginorm	16591	76		82687	81		97	100		2617095	89	
	raw	21784			102287			97			2941524		
Bacteroides	Bignorm	3356	68	100	25300	100	100	95	98	99	70206	105	112
	Diginorm	3356	68		25300	100		96	99		62882	94	
	raw	4930			25299			98			66626		
Caldi	Bignorm	50973	82	83	143346	89	91	100	100	100	573836	94	68
	Diginorm	61108	98		157479	98		100	100		839126	138	
	raw	62429			160851			100			609604		
Caulo	Bignorm	4515	69	95	20255	100	107	96	98	98	60362	86	113
	Diginorm	4729	72		18907	93		98	101		53456	76	
	raw	6562			20255			97			70161		
Chloroflexi	Bignorm	13418	102	109	79605	102	102	99	100	100	666519	95	93
	Diginorm	12305	93		78276	100		100	100		716473	102	
	raw	13218			78276			99			703171		
Crenarch	Bignorm	6538	77	91	31401	81	66	97	99	99	484354	89	95
	Diginorm	7148	84		47803	124		98	100		510256	94	
	raw	8501			38582			98			544763		
Cyanobact	Bignorm	5833	95	99	33462	98	100	99	101	100	236391	113	110
	Diginorm	5907	96		33516	98		99	101		214574	103	
	raw	6130			34300			98			209269		
E. coli	Bignorm	112393	100	100	268306	94	94	96	100	100	28966	65	65
	Diginorm	112393	100		285311	100		96	100		44465	100	
	raw	112393			285528			96			44366		
SAR324	Bignorm	135669	100	114	302443	100	100	99	100	100	4259479	98	100
	Diginorm	119529	88		302443	100		99	100		4264234	98	
	raw	136176			302442			99			4342602		

Table 5. Filter and assembly statistics for Bignorm with $Q_0 = 20$, Diginorm and the raw datasets (II)

	Min	Q1	Median	Mean	Q3	Max
<u>Diginorm mean phred score</u>						
raw mean phred score	62	66	74	74	79	89
<u>Bignorm filter time</u>						
Diginorm filter time	24	28	31	33	38	46
<u>Bignorm SPAdes time</u>						
Diginorm SPAdes time	4	08	18	26	35	88

Table 6. Quartiles for comparison of mean phred score, filter and assembly time in %.

4 Discussion

The quality parameter Q_0 that Bignorm introduces over Diginorm has shown to have a strong impact on the number of reads kept, coverage, and quality of the assembly. An upper bound of $Q_0 \leq 25$ for a reasonable Q_0 was obtained by considering the 10th percentile of the coverage (Figure 2b). With this constraint in mind, in order to have a small number of reads kept, Figure 1a suggests $18 \leq Q_0 \leq 25$. Given that N50 for E.coli starts to decline at $Q_0 = 20$ (Figure 3), we decided for $Q_0 = 20$ as the recommended value. As seen in detail in Table 4, $Q_0 = 20$ gives good assemblies for all 13 cases. The gain in speed is considerable: in terms of median we only require 31% and 18% of Diginorm’s time for filtering and assembly, respectively. This speedup generally comes at the price of a smaller genomic fraction and smaller largest contig, although those differences are relatively small.

5 Conclusions

For 13 bacteria single cell datasets, we have shown that good and fast assemblies are possible, based on only 5% of the reads in most of the cases (and on less than 10% of the reads in all of the cases). The filtering process, using our new algorithm Bignorm, also works fast and much faster than Diginorm. Like Diginorm, we use a count-min sketch for counting k -mers, so our the memory requirements are relatively small and known in advance. We provide tuning for the quality parameter Q_0 and recommend to use $Q_0 = 20$ in practice. We refrained from tuning the other parameters c_0 , c_1 that are used to define the contributions $b_0(s)$ and $b_1(s)$, as well as the N-count threshold N_0 and contribution threshold B . We expect that tuning of those parameters will help to obtain assemblies of higher quality and intend to do so in future work.




Figure 3. Statistics for the assemblies of four selected datasets.

References

- [1] Acetothermia bacterium JGI MDM2 LHC4sed-1-H19. URL: http://genome.jgi.doe.gov/AcebacLHC4se1H19_FD/AcebacLHC4se1H19_FD.info.html.
- [2] Alphaproteobacteria bacterium SCGC AC-312_D23v2. URL: http://genome.jgi.doe.gov/AlpbacA312_D23v2_FD/AlpbacA312_D23v2_FD.info.html.

- [3] *Arcobacter* sp. SCGC AAA036-D18. URL: http://genome.jgi.doe.gov/ArcspSAAA036D18_FD/ArcspSAAA036D18_FD.info.html.
- [4] *Armatimonadetes* bacterium JGI 0000077-K19. URL: http://genome.jgi.doe.gov/Armbac0000077K19_FD.
- [5] *Bacteroidetes* bacVI JGI MCM14ME016. URL: http://genome.jgi.doe.gov/BacbacMCM14ME016_FD.
- [6] *Calescamantes* bacterium JGI MDM2 SSWTFF-3-M19. URL: http://genome.jgi.doe.gov/CalbacSSWTFF3M19_FD.
- [7] *Candidatus Poribacteria* sp. WGA-4E. URL: http://genome.jgi.doe.gov/CanPorspWGA4E_FD.
- [8] *Caulobacter* bacterium JGI SC39-H11. URL: http://genome.jgi.doe.gov/CaubacJGISC39H11_FD.
- [9] *Chloroflexi* bacterium SCGC AAA257-O03. URL: http://genome.jgi.doe.gov/ChlbacSAAA257O03_FD.
- [10] *Crenarchaeota* archaeon SCGC AAA261-F05. URL: http://genome.jgi.doe.gov/CrearcSAAA261F05_FD.
- [11] *Cyanobacteria* bacterium SCGC JGI 014-E08. URL: http://genome.jgi.doe.gov/CyabacSJGI014E08_FD.
- [12] FASTX-Toolkit. URL: http://hannonlab.cshl.edu/fastx_toolkit/.
- [13] JGI Genome Portal - Home. URL: <http://genome.jgi.doe.gov>.
- [14] Single cell data sets. URL: http://bix.ucsd.edu/projects/singlecell/nbt_data.html.
- [15] Andy S. Alic, David Ruzafa, Joaquin Dopazo, and Ignacio Blanquer. Objective review of de novo stand-alone error correction methods for NGS data. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 6(2):111–146, 2016. doi:10.1002/wcms.1239.
- [16] Anton Bankevich, Sergey Nurk, Dmitry Antipov, Alexey A. Gurevich, Mikhail Dvorkin, Alexander S. Kulikov, Valery M. Lesin, Sergey I. Nikolenko, Son Pham, Andrey D. Prjibelski, Alexey V. Pyshkin, Alexander V. Sirotkin, Nikolay Vyahhi, Glenn Tesler, Max A. Alekseyev, and Pavel A. Pevzner. SPAdes: A New Genome Assembly Algorithm and Its Applications to Single-Cell Sequencing. *Journal of Computational Biology*, 19(5):455–477, Apr 2012. doi:10.1089/cmb.2012.0021.
- [17] Anthony M. Bolger, Marc Lohse, and Bjoern Usadel. Trimmomatic: A flexible trimmer for Illumina Sequence Data. *Bioinformatics*, 30(15):2114–2120, 2014. doi:10.1093/bioinformatics/btu170.
- [18] Graham Cormode and S. Muthukrishnan. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms*, 55(1):58–75, 2005. doi:10.1016/j.jalgor.2003.12.001.

- [19] Murray P. Cox, Daniel A. Peterson, and Patrick J. Biggs. SolexaQA: At-a-glance quality assessment of Illumina second-generation sequencing data. *BMC Bioinformatics*, 11(1):1–6, 2010. doi:10.1186/1471-2105-11-485.
- [20] Michael Crusoe, Greg Edverson, Jordan Fish, Adina Howe, Eric McDonald, Joshua Nahum, Kaben Nanlohy, Humberto Ortiz-Zuazaga, Jason Pell, Jared Simpson, Camille Scott, Ramakrishnan Rajaram Srinivasan, Qingpeng Zhang, and C. Titus Brown. The khmer software package: enabling efficient sequence analysis. pages 1–3, 04 2014. doi:10.6084/m9.figshare.979190.
- [21] Cristian Del Fabbro, Simone Scalabrin, Michele Morgante, and Federico M. Giorgi. An Extensive Evaluation of Read Trimming Effects on Illumina NGS Data Analysis. *PLoS ONE*, 8(12):1–13, 12 2013. doi:10.1371/journal.pone.0085024.
- [22] Martin Dietzfelbinger, Torben Hagerup, Jyrki Katajainen, and Martti Penttonen. A Reliable Randomized Algorithm for the Closest-Pair Problem. *Journal of Algorithms*, 25(1):19–51, 1997. doi:10.1006/jagm.1997.0873.
- [23] Dent Earl, Keith Bradnam, John St John, Aaron Darling, Dawei Lin, Joseph Fass, Hung On Ken Yu, Vince Buffalo, Daniel R Zerbino, Mark Diekhans, et al. Assemblathon 1: A competitive assessment of de novo short read assembly methods. *Genome research*, 21(12):2224–2241, 2011. doi:10.1101/gr.126599.111.
- [24] Sante Gnerre, Iain Maccallum, Dariusz Przybylski, Filipe J Ribeiro, Joshua N Burton, Bruce J Walker, Ted Sharpe, Giles Hall, Terrance P Shea, Sean Sykes, Aaron M Berlin, Daniel Aird, Maura Costello, Riza Daza, Louise Williams, Robert Nicol, Andreas Gnirke, Chad Nusbaum, Eric S Lander, and David B Jaffe. High-quality draft assemblies of mammalian genomes from massively parallel sequence data. *Proc Natl Acad Sci U S A*, 108(4):1513–1518, 1 2011. doi:10.1073/pnas.1017351108.
- [25] Alexey Gurevich, Vladislav Saveliev, Nikolay Vyahhi, and Glenn Tesler. QUAST: quality assessment tool for genome assemblies. *Bioinformatics*, 29(8):1072–1075, 2013. doi:10.1093/bioinformatics/btt086.
- [26] NA Joshi and JN Fass. Sickel: A sliding-window, adaptive, quality-based trimming tool for FastQ files (Version 1.33). Available at <https://github.com/najoshi/sickle>, 2011.
- [27] Rei Kajitani, Kouta Toshimoto, Hideki Noguchi, Atsushi Toyoda, Yoshitoshi Ogura, Miki Okuno, Mitsuru Yabana, Masahira Harada, Eiji Nagayasu, Haruhiko Maruyama, Yuji Kohara, Asao Fujiyama, Tetsuya Hayashi, and Takehiko Itoh. Efficient de novo assembly of highly heterozygous genomes from whole-genome shotgun short reads. *Genome Research*, pages 1384–1395, 2014. doi:10.1101/gr.170720.113.
- [28] Janine Kamke, Alexander Sczyrba, Natalia Ivanova, Patrick Schwientek, Christian Rinke, Kostas Mavromatis, Tanja Woyke, and Ute Hentschel. Single-cell genomics reveals complex

- carbohydrate degradation patterns in poribacterial symbionts of marine sponges. *The ISME journal*, 7(12):2287–2300, 2013. doi:10.1038/ismej.2013.111.
- [29] David R Kelley, Michael C Schatz, and Steven L Salzberg. Quake: quality-aware detection and correction of sequencing errors. *Genome Biol*, 11(11):1–13, 2010. doi:10.1186/gb-2010-11-11-r116.
- [30] Ben Langmead and Steven L. Salzberg. Fast gapped-read alignment with Bowtie 2. *Nat Meth*, 9(4):357–359, 4 2012. Brief Communication. doi:10.1038/nmeth.1923.
- [31] Marcel Martin. Cutadapt removes adapter sequences from high-throughput sequencing reads. *EMBnet. journal*, 17(1):10–12, 2011. doi:10.14806/ej.17.1.200.
- [32] Nicola Prezza, Cristian Del Fabbro, Francesco Vezzi, Emanuele De Paoli, and Alberto Policriti. ERNE-BS5: Aligning BS-treated Sequences by Multiple Hits on a 5-letters Alphabet. In *Proceedings of the ACM Conference on Bioinformatics, Computational Biology and Biomedicine*, BCB ’12, pages 12–19, New York, NY, USA, 2012. ACM. doi:10.1145/2382936.2382938.
- [33] Aaron R. Quinlan and Ira M. Hall. BEDTools: a flexible suite of utilities for comparing genomic features. *Bioinformatics*, 26(6):841–842, 2010. doi:10.1093/bioinformatics/btq033.
- [34] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2016. Available at <https://www.R-project.org/>, Version 3.3.0.
- [35] Robert Schmieder and Robert Edwards. Quality control and preprocessing of metagenomic datasets. *Bioinformatics*, 27(6):863–864, 2011. doi:10.1093/bioinformatics/btr026.
- [36] Linnéa Smeds and Axel Künstner. ConDeTri - A Content Dependent Read Trimmer for Illumina Data. *PLoS ONE*, 6(10):1–6, 10 2011. doi:10.1371/journal.pone.0026314.
- [37] C. Titus Brown, A. Howe, Q. Zhang, A. B. Pyrkosz, and T. H. Brom. A Reference-Free Algorithm for Computational Normalization of Shotgun Sequencing Data. *ArXiv e-prints*, pages 1–18, March 2012. arXiv:1203.4802.
- [38] Philipp Wölfel. *Über die Komplexität der Multiplikation in eingeschränkten Branchingprogrammmodellen*. PhD thesis, Universität Dortmund, Fachbereich Informatik, 2003.
- [39] Qingpeng Zhang, Jason Pell, Rosangela Canino-Koning, Adina Chuang Howe, and C. Titus Brown. These Are Not the K-mers You Are Looking For: Efficient Online K-mer Counting Using a Probabilistic Data Structure. *PLoS ONE*, 9(7):1–13, 07 2014. doi:10.1371/journal.pone.0101271.